

Symbolic Dynamics for Discrete Adaptive Games

Cosma Rohilla Shalizi^{1,2,*} and David J. Albers^{3,2,†}

¹*Center for the Study of Complex Systems, University of Michigan, Ann Arbor, MI 48109*

²*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501*

³*Physics Department, University of Wisconsin, Madison, WI 53706*

(Dated: February 1, 2008)

We use symbolic dynamics to study discrete adaptive games, such as the minority game and the El Farol Bar problem. We show that no such game can have deterministic chaos. We put upper bounds on the statistical complexity and period of these games; the former is at most linear in the number of agents and the size of their memories. We extend our results to cases where the players have infinite-duration memory (they are still non-chaotic) and to cases where there is “noise” in the play (leaving the complexity unchanged or even reduced). We conclude with a mechanism that can reconcile our findings with the phenomenology, and reflections on the merits of simple models of mutual adaptation.

PACS numbers: 87.23.Ge, 89.65.Gh, 89.75.-k, 05.45.-a

Keywords: Minority Game, El Farol Bar problem, discrete adaptive games, bounded rationality, symbolic dynamics, statistical complexity, exponentially-long transients

I. DISCRETE ADAPTIVE GAMES

A. Introduction

In the last few years, physicists have become very interested in discrete games. In these models, a finite number of players or agents pick from a finite set of actions (possibly varying from agent to agent), at discrete times, and receive rewards that depend on their own and the other agents’ actions. Each agent chooses according to a rule called its *strategy*; if agents have more than one strategy available, they use a meta-strategy to choose among them.

As described, the setting fits classical game theory [1, 2], with its assumption of so-called substantive rationality — the agents have rational, accurate expectations, enjoy common knowledge of the structure of the game, and act so as to maximize their reward. Casual observation suggests this is not a good model of human behavior, and indeed experimental psychology [3] and experimental economics [4] both reject it decisively. Failing substantive rationality, one turns to some form of *bounded rationality* [5] — agents still pursue goals, but are limited by finite knowledge and finite cognitive capacities, and so cannot, in general, maximize anything. Bounded rationality is the subject of a rich literature in the social sciences, which encompasses both abstract models [6], some of which exploit familiar methods of statistical physics [7], and a rapidly-growing experimental literature, which includes its own array of models [8, 9, 10, 11]. While the econophysical literature almost exclusively considers boundedly-rational agents, it never builds upon other work on bounded rationality. For instance, one of the best-supported meta-strategies is that

of *satisficing*, in which an agent searches over its options and takes the first one it finds with a satisfactory reward [12, 13, 14]. We are unable to find a single paper which employs this notion. Econophysical strategies are all, so far as we know, of the following form. Each agent remembers a simple summary statistic (e.g., the best action to have taken) for each of the last m turns of the game. For each such *history*, a strategy specifies an action to take. Rather than maximizing, or even satisficing, they employ adaptive meta-strategies tending to select strategies that did well in the recent past.

These caricatures or toy models have been studied since the 1980s, because, it is said, they (1) shed light on real learning and mutual adaptation, and (2) show surprisingly complex behavior. In this paper, we examine two forms this complexity is held to take — deterministic chaos, which is a sort of dynamical complexity, and formally measured complexity, specifically the amount of information embedded in the system. The rest of this section describes the most popular discrete adaptive games, the minority game and the El Farol Bar problem. The next section introduces the tools from symbolic dynamics, particularly subshifts of finite type, which we will use to pry open the mechanisms of these games. We show that the games are non-chaotic, and put linear upper bounds on their complexities. We discuss how to extend our results to games with either memories of unlimited length, or noisy play. We then consider what you must do to a game to make it chaotic, if you want it to be chaotic. In closing, we suggest a way to reconcile our findings with common observations on these games.

B. Examples: El Farol and the Minority Game

The best-loved game in econophysics is the *minority game* of Challet and Zhang [15], the subject of a rapidly growing literature (see

*Electronic address: cshalizi@umich.edu

†Electronic address: albers@cow.physics.wisc.edu

<http://www.unifr.ch/econophysics/minority/>).

There are N agents, each of which can make only one of two moves: play “0” or play “1”. If most agents play 0, then those who play 1 win, and vice-versa. (N is always odd.) Each player remembers whether 0 or 1 won on each of the last m turns. Strategies, again, are rules mapping histories to actions; each agent has s of them. In effect, strategies must predict what most players will do, and then do the opposite. Agents keep track of how many times each of their strategies would have won, and use the strategy with the best record over the time it can remember. (Note that psychologists have studied human performance in just such discrete prediction situations since the 1950s, without finding anything resembling this learning mechanism [16].) Ties between strategies are broken randomly or pseudo-randomly, say by assigning strategies a random order at the beginning of the game.

Traditionally, the minority game has been studied either using the brute empiricism of simulations, or with statistical-mechanical methods, particularly those employing the thermodynamic limit. Recently, Jefferies, Johnson et al. have introduced methods owing more to dynamical systems, e.g., de Bruijn graphs over possible sequences of winning outcomes [17, 18]. Our approach here is related to theirs, but is even more deterministic, and more microscopic.

The minority game descends from W. B. Arthur’s El Farol Bar Problem [19]. There are 100 people in Santa Fe who go out to bars; El Farol is the only good bar. Normally, those who go there are better off than those who stay home and watch the stars. But El Farol is small, and if more than 40 people go, crowding makes the bar-goers worse off than stay-at-homes. Everyone knows how many people went to the bar on each of the last m nights. Strategies, again, map histories to actions — staying home or going to the bar.

Rather than statistical mechanics, the traditional approaches to the El Farol problem are those of evolutionary adaptation and inductive behavior. Many have sought to show that either the population of agents co-evolves a set of mutually-tolerable strategies, or that a ceaseless evolutionary arms race forces everyone to “run as fast as they can just to stay in place” [2, 19].

II. SYMBOLIC DYNAMICS, SHIFTS

Symbolic dynamics is one of the basic analytical methods of nonlinear dynamics and complex systems, and as such its elements are discussed in standard introductory textbooks on those subjects (e.g., Refs. [20, 21, 22, 23, 24, 25]). Strangely, however, it does not seem to have been applied to discrete adaptive games before. We therefore give a few words of motivation, before presenting some basic notions and results. Readers whose curiosities are aroused will find Refs. [24] and [25] particularly nice introductions, emphasizing uses in nonlinear dynamics and complexity, respectively. Refs. [26] and [27] more

advanced treatments.

A. Why Symbolic Dynamics?

The essence of symbolic dynamics is substituting strings of discrete characters or symbols for trajectories in an original continuous state space. Motion in the original space corresponds simply to shifting along the string. This idea goes back to work by Morse and others in the 1930s. One of its first great successes came in a paper by Smale [28] which introduced his famous “horseshoe”, and established the geometry of what we now call chaotic dynamics. Smale first constructed a differentiable, invertible map f on a two-dimensional manifold M , which, while it has a very simple functional form, has almost impossibly complicated geometric behavior. He then connected it to the full binary shift map σ on infinite sequences of 0s and 1s — that is, the map which simply substitutes the symbol at position $i + 1$ for that at position i . The dynamics of σ are very easy to grasp. In particular, it is easy to see that it has a countable infinity of periodic points of arbitrarily high period, an uncountable infinity of non-periodic points, and a dense orbit. (Recall that rational numbers, which are countable, have repeating binary representations, while irrational numbers, which are uncountable, never repeat.) The great trick, then, was to show that the horseshoe map f and the shift map σ were topologically conjugate. That is, he found a continuous function ϕ from a subset of the manifold M to the sequence space Σ which was one-to-one and onto, and which commuted with the two maps. (Symbolically, $\phi \circ f = \sigma \circ \phi$.) This meant that the dynamical properties of the horseshoe map could be read off directly from those of the shift map.

The attraction of symbolic dynamics should now be clear: we can, so to speak, factor the dynamics of interest into a clean, simple shift map (σ), and a messier, but less material, conjugacy map (ϕ). If, for instance, we could show that the shift map and the minority game are topological conjugate, or discern the mechanisms that prevent them from being conjugated, the dynamics of the minority game itself would be much more plain to us.

B. Basic Notions

We now will quote the relevant standard results without proofs, which may be found in e.g. Ref. [26]. Begin with a finite or countable *alphabet* X of *symbols*. From these we make *strings* or *words* over X , which are *symbol sequences*, $x_0x_1x_2\dots x_n$. Our *shift space* is a set of semi-infinite words, $x = x_0x_1x_2\dots x_i\dots$ (i.e. we will be concerned with sequences on only one side of the decimal point). The dynamics are given by a *shift operator* T : $Tx = x_1x_2\dots$. That is, $(Tx)_i = x_{i+1}$. The *orbit* of a point is its set of future iterates, $\{x, Tx, T^2x, \dots\}$.

It is useful, and perhaps comforting, to invoke linear operators at this point. Let A be a square 0-1 matrix, with as many rows as X has symbols. Define a shift space X_A where x is allowed iff $\forall i, A_{(x_i)(x_{i+1})} = 1$. Then X_A is a *subshift of finite type* and the matrix A corresponds to the shift operator T . We can also represent X_A with a directed graph G_A , where there is one node for each symbol, and an edge from i to j iff $A_{ij} = 1$. Then $x \in X_A$ iff it is a path through G_A . The transition matrix is *one-many* if there are symbols i, j, k such that $A_{ij} = 1$ and $A_{ik} = 1$ but $j \neq k$. A is *many-one* if there are i, j, k such that $A_{ik} = 1$ and $A_{jk} = 1$ but $i \neq j$. And A is *one-one* if it is neither many-one nor one-many.

We now have a space X_A of sequences, and some of the properties of this space will be useful later on. First, the standard metric between two sequences (points) x and y in the shift space is

$$d(x, y) = \sum_{i=0}^{\infty} \frac{1 - \delta(x_i, y_i)}{2^i}, \quad (1)$$

where $\delta(a, b)$ is the Kronecker delta. Thus if $x_i = y_i$ for all $i \leq M$, then $d(x, y) \leq 2^{-M}$. Conversely, if $d(x, y) < 2^{-M+1}$, then $x_i = y_i$ for all $i \leq M$. From this it follows that the sequence space, equipped with the standard metric, is compact, totally disconnected and perfect (i.e., it is closed, and every point is a limit point).

C. Entropies and Lyapunov Exponents

It is now useful to deal with notions defining growth rates of a given space and spreading rates along orbits. This will be achieved via standard entropies and Lyapunov exponents borrowed from ergodic theory. In general, not all words are allowed, i.e., not every semi-infinite word is in the shift space. Likewise, many of the sequences are never seen due to the particulars of the shift operator. If $W(L)$ is the number of allowed words of length L , the *topological entropy* is

$$h = \lim_{L \rightarrow \infty} \frac{1}{L} \log W(L). \quad (2)$$

That is, h is the asymptotic exponential rate of growth in the number of allowed words. It measures the degree of exponential spreading in the dynamics. Normally, of course, that is measured by the Lyapunov exponents. Assuming those can be defined, the largest Lyapunov exponent $\lambda_1 \leq h$. (For specific conditions under which the Lyapunov exponents exist, see Refs. [29, 30, 31].)

The topological entropy gives equal weight to every *possible* word, no matter how improbable. If we wish to give more weight to the more probable strings, it is natural to start with the entropy over words of length L ,

$$H(L) \equiv \langle -\log \Pr(x_1^L) \rangle \quad (3)$$

where $\langle \rangle$ denotes expectation. This is the mean amount of information, in bits, required to specify a word of length

L generated by the dynamics. We then define the *metric entropy rate* or *Kolmogorov-Sinai entropy* as

$$h_\mu \equiv \lim_{L \rightarrow \infty} \frac{H[L]}{L}. \quad (4)$$

In addition to being an important invariant of the dynamical system, the metric entropy rate can be thought of as the rate of *information production* — the rate at which new information is produced by the dynamics. Clearly, $h \geq h_\mu$ [32]. Under certain conditions on f [30, 33, 34], h_μ is equal to the sum of the positive Lyapunov exponents; more generally it is a lower bound on their sum. Thus, a positive h_μ guarantees the existence of at least one positive Lyapunov exponent. The proofs of these assertions are somewhat complicated results from ergodic theory, well-summarized in Ref. [32].

D. Chaos

We adopt the standard topological definition of chaos, due to Devaney [35]. A compact set Λ , invariant under T , is *chaotic* if T is topological transitive on Λ , and its periodic orbits are dense in Λ . *Topological transitivity* means that the trajectory of any open set in Λ eventually visits every other open set in Λ — we can get from one part of the region to any other. Together with the requirement of dense periodic orbits, this implies sensitive dependence on initial conditions. One can imagine the chaotic trajectories wandering between (but never on) the dense periodic orbits. Note that the periodic orbits, while dense, are not stable or attracting [36]. In fact, sensitive dependence on initial conditions, which we have just established, in turn implies exponential separation of nearby orbits, at least under fairly weak conditions, and hence a positive Lyapunov exponent and positive entropy rates. These, accordingly, constitute the numerical or experimental signatures of chaos.

For subshifts of finite type, all of this can be related to the matrix A (or the corresponding graph) very directly. X_A is *transitive* if it has dense orbits. (N.B., not necessarily dense *periodic* orbits.) X_A is *irreducible* if for every pair of indices i and j , there is some $l > 0$ for which $\{A^l\}_{ij} > 0$. (This is the same as G_A being strongly connected.) Now, if X_A is irreducible, then either it is transitive or it consists of a single periodic orbit. Likewise, if X_A is transitive, then A is irreducible.

III. THE MINORITY GAME AS A SHIFT MAP

To do one iteration of the minority game, we need to know two things. (1) The history of the last m rounds — which action won on each of them. (2) The “gains” of all the strategies possessed by the agents — the number of times each strategy predicted the winning action in the last m rounds. These facts fix the action of each agent, and so the winner of the round. That is, the next

state of the game is a deterministic function of its present state. (In what follows, “state” will always mean the global state of the game, never the state of an individual agent.)

Since each state variable takes only finitely many values, we can map the state as a whole into a finite alphabet. We order the states and their alphabetic symbols however we like. Then we build a square matrix A , where $A_{ij} = 1$ if j follows i , and $= 0$ otherwise. X_A defines a subshift of finite type over states of the minority game, and every run of the minority game is an allowed orbit of this subshift. We could write down A in some detail, but shan’t.

Suppose memory for gains extends back indefinitely. Then the state comes from a countably infinite alphabet, rather than a finite one, but the subshift is still deterministic (many-one).

In the finite-memory case, the lack of chaos is elementary once we realize that there are only finitely many states; the full apparatus of symbolic dynamics is overkill. There are only finitely many symbols, so every string will eventually repeat a symbol. But each symbol has a unique successor, so after the first repetition, the string just repeats a single word. Hence every orbit is eventually periodic.

The topological entropy rate is zero. The first symbol of a string fixes all subsequent ones, so $\forall L$, $W(L) = W(1)$, and $h = 0$ by Eq. 2.

IV. COMPLEXITY IN THE MINORITY GAME

Before we can say quantitative things about the complexity of a model, we must pick a complexity measure, of which there are many [25]. When people have discussed the complexity of discrete games at all, they have picked the Kolmogorov complexity or “algorithmic information content” [37]. Kolmogorov introduced this concept as part of a definition of “randomness” relying on purely combinatorial notions, part of a long-standing program of research in the foundations of probability [38]. The Kolmogorov complexity $K(w)$ of a finite string w is the length of the shortest computer program which will produce that string as its output and then stop [66]. Because any string can be produced by a program like `print ‘w’`, the complexity of any string cannot be more than its length plus a constant (the length of the “print” instruction). On the other hand, some strings can be produced by programs which are much shorter than themselves, e.g., a string of a million 0s can be produced by iterating the command `print ‘0’` a million times. A random string, roughly speaking, is one which cannot be generated by any program significantly shorter than the string itself. More exactly, if x_1^L are the first L symbols of an infinite word, and $K(x_1^L)/n \rightarrow 1$, then all the properties which are true of a sequence of independent, identically distributed random variables are true of the symbols in the infinite word. For deep reasons connected with the

existence of uncomputable numbers, there can be no general procedure for computing the Kolmogorov complexity of arbitrary strings, nor are there any known procedures which give well-controlled approximations [37, 39]. However, if the symbols are generated by an ergodic process, there is an easy expression for the expected rate of growth in the complexity [40]: $\langle K(x_1^L) \rangle / L \rightarrow h_\mu$ as $L \rightarrow \infty$. This is another illustration of the fact that h_μ measures the rate of information production.

While, as we have said, there is no real way to compute the Kolmogorov complexity, several papers in the physics literature have attempted to calculate it for games [41, 42, 43] and even for real markets [44, 45], with the claim [45] that the high complexity they obtain shows that financial time series contain much valuable information. *Pace* these authors, had they actually been able to calculate the algorithmic information content, a high value of it would indicate that the systems they considered were basically random, because this is just what Kolmogorov complexity *means*.

Assuming one wants a notion of “complex” which is not the same as “random”, what should a good complexity measure be like? As mentioned, many measures have been proposed [25], but there is some consensus about what features a good measure should have. It should be low for both highly ordered and highly random systems; it should be calculable; and it should have a clear dynamical meaning, rather than just being a number for its own sake [46]. To our knowledge, only one measure satisfies all these criteria, the *statistical complexity* of Crutchfield and Young [47], a fully formalized and operational version of Grassberger’s [48] conjectural “true measure complexity”. It is the amount of information about the past of the system needed for maximally accurate prediction of its future [49].

We get the statistical complexity by finding the *causal states* of the system: each causal state is a distinct distribution for future events, conditional on the past of the system. Abstractly, we find the conditional distribution $\Pr(x_1^\infty | x_{-\infty}^0)$ for each past $x_{-\infty}^0$, and group pasts with the same distribution into states. The probability of a state is then the sum of the probabilities of all the pasts assigned to it, and the distribution conditional on the state is just the same as the distribution conditional on the individual pasts. (In practice, we use a number of recursive tricks to approximate the causal states from finite data, while only looking at the conditional distribution of the next symbol [50].) The statistical complexity C_μ is simply the entropy (uncertainty in bits) of the causal states $= -\sum_i p_i \log_2 p_i$, where p_i is the probability of the i^{th} causal state. This is maximized when all causal states are equally probable. In the case of a periodic process, each phase is a causal state, and each state is equally probable, so $C_\mu = \log_2 P$. [67]

We have just seen that the minority game has only periodic orbits, so to get C_μ , we just need the period. Start with just one strategy per agent. After the system has seen every possible history, it must repeat itself, so

$P \leq 2^m$. Now give each agent several strategies. The number of ways agents can deploy strategies is $S \leq s^N$. When we've gone through every possible history, the agents could be using different strategies than before. But once we've seen all deployments of strategies, we must be back where we started. Thus $P \leq S2^m \leq s^N 2^m$. Even these bad estimates give us

$$C_\mu \leq m + N \log_2 s \quad (5)$$

The complexity grows linearly with memory and population *at most*. We get this growth rate only with complete (combinatorial) independence — an agent's strategy selection is constrained neither by the selections of other agents, nor even by the global history!

In a number of recent papers, Savit, Parunak and collaborators [51, 52, 53] have considered a complexity measure closely related to the statistical complexity (they do not give it a name). In essence, following [54], they treat two pasts as belong to the same state if and only if they lead to the same *unique* next symbol, whereas statistical complexity requires only that they lead to the same *distribution* of future symbols. The two complexities coincide for periodic processes. In the presence of noise, however, the method of Savit et al. will be forced to rely on longer and longer histories, until it finds ones that look deterministic simply because of undersampling, and their complexity will diverge logarithmically in the size of their data set [55]. The statistical complexity, however, is well-behaved in the presence of noise (see Section VI below).

V. INFINITE MEMORY FOR GAINS

People sometimes play the minority game with an infinite memory for gains — agents remember the number of times each strategy has won since the beginning of the game, rather than just in the last m turns. Then the alphabet is countably infinite. But what matters are the difference in strategies' gains, and if those are bounded, the finite-state arguments take over. Even when differences in gains are unbounded, the shift map is many-one, so there is just a single 1 in each row of the transition matrix T .

The topological entropy rate of a countable-state shift is [26, Observation 7.2.10]

$$h = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_j (T^n)_{ij}, \quad (6)$$

independent of i . From many-oneness, the sum in Eq. 6 is always exactly 1, so $h = 0$.

Let us consider also the question of exponential spreading.

The distance between points in the standard metric is at most 2. $d(x, y) = 2$ iff $x_i \neq y_i$ for all i . In that case, $d(T^n x, T^n y) = 2$ for all n , so maximally-separated points stay maximally separated. On the other hand, if

the distance between x and y is less than 2, then at some position, k , $x_k = y_k$. By the many-oneness of the shift, if $j > k$, $x_j = y_j$. Therefore, the distance between the iterates of x and y shrinks until it hits 0. Distance never increases; spreading of any sort, never mind exponential spreading, is impossible.

The distance between any two distinct periodic points is always maximal; that is, given two periodic points x and y , if $x \neq y$, then $x_i \neq y_i$ for all i . To see this, consider first the case of two points on the same periodic cycle, of period P . For some $0 < L < P$, $T^L(x) = y$. Suppose that, for some k , $x_k = y_k$. Then $T^k x = T^k y = T^{k+L} x$, so $T^k x$ is periodic with period $L < P$, which is absurd. Now consider points on two distinct periodic cycles. Suppose there were a k such that $x_k = y_k$. Then $T^k(x) = T^k(y)$. Hence there is a point, $T^k(x)$, belonging both to the x cycle and the y cycle. But this is absurd, because the cycles are distinct.

Suppose periodic points were dense. Then for any point x and any distance $\epsilon > 0$, there would be a periodic point $y \neq x$ such that $d(x, y) \leq \epsilon$. But if x is itself periodic, then we have just seen that $d(x, y) = 2$ no matter what periodic point y we chose. Therefore periodic points are not dense. Chaotic maps have dense periodic points, so this shift map isn't one.

VI. NOISE

Traditionally, the minority game is played somewhat noisily — either players chose strategies with some degree of randomness, or the move they make, once they have chosen a strategy, is somewhat random. This makes the game amenable to statistical-mechanical treatment, a point which has been thoroughly explored elsewhere. Here, we will see that it can make the entropy rate positive, but does not increase the complexity.

The usual types of noise imposed in discrete games change them from deterministic maps into Markov chains. This means we can represent the state of the game at one time as a function of the state at the previous time, and a sequence of stationary, independent noise variables, i.e., $x_{t+1} = f(x_t, \eta_t)$, where $\{\eta_t\}$ are IID random variables which are also independent of the x_t . The Markov property of the sequence $\{x_t\}$ also means that its entropy rate (metric or topological) is just the entropy of x_{t+1} conditional on x_t [26, 39]. That is, $h_\mu = H[x_{t+1}|x_t]$, and similarly for h . Now, applying some basic results of information theory [39, ch. 2],

$$\begin{aligned} H[x_{t+1}|x_t] &= H[f(x_t, \eta_t)|x_t] \\ &\leq H[x_t, \eta_t|x_t] \\ &= H[\eta_t|x_t] = H[\eta_t] \end{aligned}$$

Or, in words: because the only source of randomness is the noise, the entropy rate of the states must be no more than the entropy of the noise. A parallel argument holds for the topological entropy. Nothing in the process *amplifies* the noise.

Now let us consider the statistical complexity. Recall that a causal state is a distinct distribution of the system's future behavior, conditional on its past. For noise of the sort described, the present state of the game determines the *distribution* of future symbols, independent of any other feature of the past of the system. So the present symbol fixes the causal state. The noise could make some causal states more likely than others (reducing the complexity), or could even lead to two symbols having the same distribution over future orbits, merging them into a single causal state (again lowering the complexity). Adding noise simplifies matters, if it does anything at all.

VII. CONCLUSIONS

We analyzed the global behavior of discrete adaptive games, using symbolic dynamics. These games are not chaotic, and the most popular ones are actually of low complexity, at most linear in the number of agents.

The crucial obstacles to chaos were that the state of the minority game can be summarized in a discrete alphabet, and that the evolution of those states is many-one. Our results will apply, *mutatis mutandis*, to any deterministic adaptive game of which this is true, for instance, the El Farol problem, and various discrete market models. To repeat, this is not an artifact of discretizing continuous models, but a feature of deliberately discrete ones. Nothing depends on the meta-strategy; our results hold with learning, imitation, evolution, etc., so long one implements them non-anticipatively.

A. Exponential Spreading versus Exponential Transients

Anyone who looks at a simulation of the minority game sees an intricate, hard-to-predict process — we certainly do. Scientists trained in nonlinear dynamics naturally look to exponential divergence of orbits to explain these observations. We have seen that there is no such divergence in the standard, noiseless version of these games; there are three ways to get it.

- *Hidden information.* If only some agents have access to a (changing) outside source of information, then there can be apparent transitivity in the game, and so deterministic chaos. A neat example is the Brock-Hommes asset-trading model [56, 57, 58], where agents get simple adaptive strategies for free, but must pay for rational expectations of future prices.
- *Add noise.* As we've seen (Sec. VI), this certainly allows exponential divergence, but in a trivial way, and at a cost in complexity.

- *Coupling to chaos.* Break ties using a chaotic shift map, or something similar. Often this will make the game chaotic.

Only the first of these mechanisms has any *economic* relevance. In real markets, information is expensive, but transparency is prohibitively expensive, and so hidden information is ubiquitous. Exploring the consequences of this is an important problem of great economic interest; indeed, the problem is common across almost all social situations. The other methods, by contrast, are “unphysical,” lacking in substantive content. Of course, hidden information cannot explain the phenomenology of games in which it is absent, but it is no less interesting on that account.

What however *does* account for that phenomenology? We suggest that the answer is exponentially fast divergence, but rather exponentially slow convergence. For the minority game, it is numerically well established that relaxation times grow exponentially with m [59]. What is less appreciated is that, in high-dimensional dynamical systems, the approach to the attractor can be so slow that it is never reached in simulations, at least not within the lifespan of ordinary experimenters (but see [60]). In systems formed by coupling smaller dynamical systems, such as the agents in our games, the duration of the transient epoch often grows exponentially with the number of agents [61]. The attractors and their invariant distributions are irrelevant; what matters is the structure of the attractor basin [62] and the quasi-stationary transient distribution it produces [63]. Once recognized, these long transients are actually a *feature* of discrete adaptive games, since the social systems they are supposed to model are essentially never in an invariant state, but always “in transition” [68]. The transients are, however, bad news for methods based on either ergodic or thermodynamic limits. Quantities like the entropy rate and the Lyapunov exponents, which are rigorously defined only in the limit, lose much of their interest.

B. Directions for Future Work

The motive for studying models like the minority game should *not* be that they are complex systems in some formal sense. In any area of modeling, formal constructs acquire interest from their ability to help resolve problems of physical relevance — that is, here, social relevance. We have argued that such constructs as algorithmic complexities and stationary and thermodynamic limits have not demonstrated such relevance *here*, and are unlikely to do so in the future. By contrast, the distinctively transient dynamics of such systems, especially their basin structure, hold great promise for understanding the origins of complex behavior we see in these games. The hope of econophysics is that these models portray, in stylized miniature, key aspects of real mutual adaptation and social interaction, and that studying the portrait will help us understand the original. This is not an outlandish

hope, since profound insights have come from studying the equally-stylized Prisoners' Dilemma [64], and a host of other, less famous models [2, 5, 65] have illuminated the collective life of political, boundedly-rational animals. We hope that the discrete adaptive games of econophysics will take their place in this set. The contribution of econophysics, it seems to us, be to assimilate the rich body of empirical findings and models of individual choice and learning, and say something distinctive and illuminating about how large numbers of such boundedly-rational agents interact.

Acknowledgments

We thank Sam Bowles, Sven Brueckner, Jim Crutchfield, Dee Dechert, Richard Gonzalez, Cars Hommes,

Matt Jones, Cris Moore, Scott Page, Van Parunak, Robert Savit, Kris Shalizi, Clint Sprott and Libby Wood for valuable conversation and suggestions. CRS also thanks Kris Shalizi for being perfect. Our work at SFI was supported by the Dynamics of Learning program, under DARPA cooperative agreement F30602-00-2-0583, and CRS's work at Ann Arbor by a grant from the James S. McDonnell Foundation.

-
- [1] J. von Neumann and O. Morgenstern, *Theory of Games and Economic Behavior* (Princeton University Press, Princeton, 1944).
 - [2] H. Gintis, *Game Theory Evolving: A Problem-Centered Introduction to Modeling Strategic Interaction* (Princeton University Press, Princeton, 2000).
 - [3] D. Kahneman, P. Slovic, and A. Tversky, eds., *Judgment Under Uncertainty* (Cambridge University Press, Cambridge, England, 1982).
 - [4] C. Camerer, in *The Handbook of Experimental Economics*, edited by J. H. Kagel and A. E. Roth (Princeton University Press, Princeton, New Jersey, 1995), pp. 587–703.
 - [5] H. A. Simon, *The Sciences of the Artificial* (MIT Press, Cambridge, Massachusetts, 1996), 3rd ed., first edition 1969.
 - [6] A. Rubinstein, *Modeling Bounded Rationality* (MIT Press, Cambridge, Massachusetts, 1999).
 - [7] H. P. Young, *Individual Strategy and Social Structure: An Evolutionary Theory of Institutions* (Princeton University Press, Princeton, 1998).
 - [8] G. Klein, *Source of Power: How People Make Decisions* (MIT Press, Cambridge, Massachusetts, 1998).
 - [9] C. E. Zsombok and G. Klein, eds., *Naturalistic Decision Making*, Expertise: Research and Applications (Lawrence Erlbaum Associates, Mahwah, New Jersey, 1997).
 - [10] G. Gigerenzer, P. M. Todd, and the ABC Research Group, *Simple Heuristics That Make Us Smart* (Oxford University Press, Oxford, 1999).
 - [11] G. Gigerenzer and R. Selten, eds., *Bounded Rationality: The Adaptive Toolbox* (MIT Press, Cambridge, Massachusetts, 2001).
 - [12] H. A. Simon, *Quarterly Journal of Economics* **69**, 99 (1955), reprinted in [14].
 - [13] H. A. Simon, *Psychological Review* **63**, 129 (1956), reprinted in [14].
 - [14] H. A. Simon, *Models of Bounded Rationality* (MIT Press, Cambridge, Massachusetts, 1982).
 - [15] D. Challet and Y.-C. Zhang, *Physica A* **246**, 407 (1997), URL arxiv.org/abs/cond-mat/9708006.
 - [16] J. L. Myers, in *Approaches to Human Learning and Motivation*, edited by W. K. Estes (Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1976), vol. 3 of *Handbook of Learning and Cognitive Processes*, pp. 171–205.
 - [17] P. Jefferies, M. L. Hart, and N. F. Johnson, *Physical Review E* **65**, 016105 (2002), URL arxiv.org/abs/cond-mat/0103259.
 - [18] P. Jefferies, D. Lamper, and N. F. Johnson, *Anatomy of extreme events in a complex adaptive system* (2002), URL arxiv.org/abs/cond-mat/0201540.
 - [19] W. B. Arthur, *Increasing Returns and Path Dependence in the Economy* (University of Michigan Press, Ann Arbor, 1994).
 - [20] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems and Bifurcations of Vector Fields* (Springer-Verlag, New York, 1983).
 - [21] R. L. Devaney, *A First Course in Chaotic Dynamical Systems: Theory and Experiment* (Addison-Wesley, Reading, Mass., 1992).
 - [22] K. T. Alligood, T. D. Sauer, and J. A. Yorke, *Chaos: An Introduction to Dynamical Systems* (Springer-Verlag, Berlin, 1996).
 - [23] R. C. Hilborn, *Chaos and Nonlinear Dynamics: An Introduction for Scientists and Engineers* (Oxford University Press, Oxford, 1994).
 - [24] S. Wiggins, *Introduction to Applied Nonlinear Dynamical Systems and Chaos* (Springer Verlag, Berlin, 1990).
 - [25] R. Badii and A. Politi, *Complexity: Hierarchical Structures and Scaling in Physics* (Cambridge University Press, Cambridge, 1997).
 - [26] B. P. Kitchens, *Symbolic Dynamics: One-sided, Two-sided and Countable State Markov Shifts* (Springer-Verlag, Berlin, 1998).
 - [27] B.-l. Hao and W.-M. Zheng, *Applied Symbolic Dynamics and Chaos* (World Scientific, Singapore, 1998).
 - [28] S. Smale, in *Differential and Combinatorial Topology: A Symposium in Honor of Marston Morse*, edited by S. S. Cairns (Princeton University Press, Princeton, New Jersey, 1965), pp. 63–80.
 - [29] D. Ruelle, *Annals of Mathematics* **115**, 243 (1982).
 - [30] Y. Pesin, *Russian Mathematical Surveys* **32**, 55 (1977).

- [31] A. B. Katok, Publications Mathematiques de l'Institut des Hautes Etudes Scientifiques **51**, 137 (1980).
- [32] A. B. Katok and B. Hasselblatt, *Introduction to the Modern Theory of Dynamical Systems* (Cambridge University Press, Cambridge, England, 1995).
- [33] D. Ruelle, Boletim da Sociedade Brasileira de Matematica **9**, 83 (1978).
- [34] G. A. Margulis, Functional Analysis and Applications **3**, 335 (1969).
- [35] J. Banks, J. Brooks, G. Cairns, G. Davis, and P. Stacey, American Mathematical Monthly **99**, 332 (1992).
- [36] S. Smale, Bulletin of the American Mathematical Society **73**, 747 (1967).
- [37] M. Li and P. M. B. Vitanyi, *An Introduction to Kolmogorov Complexity and Its Applications* (Springer-Verlag, New York, 1993).
- [38] J. von Plato, *Creating Modern Probability: Its Mathematics, Physics and Philosophy in Historical Perspective* (Cambridge University Press, Cambridge, England, 1994).
- [39] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (Wiley, New York, 1991).
- [40] A. A. Brudno, Uspekhi Matematicheskii Nauk **33**, 207 (1978).
- [41] R. Mansilla, Complex Systems **11**, 387 (1997).
- [42] R. Mansilla, Physical Review E **62**, 4553 (2000), URL arxiv.org/abs/cond-mat/9906017.
- [43] R. Mansilla, Physica A **284**, 478 (2000), URL arxiv.org/abs/cond-mat/0002331.
- [44] R. Mansilla, Physica A **301**, 483 (2001), URL arxiv.org/abs/cond-mat/0104472.
- [45] R. N. Mantegna and H. E. Stanley, *An Introduction to Econophysics: Correlations and Complexity in Finance* (Cambridge University Press, Cambridge, England, 2000).
- [46] D. P. Feldman and J. P. Crutchfield, Physics Letters A **238**, 244 (1998).
- [47] J. P. Crutchfield and K. Young, Physical Review Letters **63**, 105 (1989).
- [48] P. Grassberger, International Journal of Theoretical Physics **25**, 907 (1986).
- [49] C. R. Shalizi and J. P. Crutchfield, Journal of Statistical Physics **104**, 817 (2001), URL arxiv.org/abs/cond-mat/9907176.
- [50] C. R. Shalizi, K. L. Shalizi, and J. P. Crutchfield, Journal of Machine Learning Research **submitted** (2002), URL arxiv.org/abs/cs.LG/0210025.
- [51] R. Savit, S. A. Brueckner, H. V. D. Parunak, and J. Sauter, Tech. Rep., Altarum (2002), URL www.erim.org/~vparunak/RAGpaper.pdf.
- [52] R. Savit, S. A. Brueckner, H. V. D. Parunak, and J. Sauter, Physics Letters A **submitted** (2003), URL arxiv.org/abs/nlin.A0/0302053.
- [53] H. V. D. Parunak, S. A. Brueckner, J. Sauter, and R. Savit, in *Autonomous Agents and Multi-Agent Systems 2002* (2002), URL www.erim.org/~vparunak/AAMAS02Effort.pdf.
- [54] R. Savit and M. Green, Physica D **50**, 95 (1991).
- [55] K. Marton and P. C. Shields, Annals of Probability **23**, 960 (1994), see also an important Correction, *Annals of Probability*, **24** (1996): 541–545.
- [56] W. A. Brock and C. H. Hommes, Journal of Economic Dynamics and Control **22**, 1235 (1998).
- [57] W. A. Brock and C. H. Hommes, Econometrica **65**, 1059 (1997).
- [58] W. A. Brock and C. H. Hommes, in *The Theory of Markets*, edited by P. Herings, A. Talman, and G. van der Laan (North Holland, for the Royal Netherlands Academy of Arts and Sciences, Amsterdam, 1999), pp. 109–137.
- [59] M. Marsili and D. Challet, Physical Review E **64**, 056138 (2001), URL arxiv.org/abs/cond-mat/0102257.
- [60] F. J. Dyson, Reviews of Modern Physics **51**, 447 (1979).
- [61] J. P. Crutchfield and K. Kaneko, Physical Review Letters **60**, 2715 (1988).
- [62] A. Wuensche and M. Lesser, *The Global Dynamics of Cellular Automata: An Atlas of Basin of Attraction Fields of One-Dimensional Cellular Automata* (Addison-Wesley, Reading, Massachusetts, 1992).
- [63] J. P. Crutchfield, Nuclear Physics B, Proceedings Supplements **5A**, 287 (1988).
- [64] R. Axelrod, *The Evolution of Cooperation* (Basic Books, New York, 1984).
- [65] T. C. Schelling, *Micromotives and Macrobehavior* (W. W. Norton, New York, 1978).
- [66] We will ignore some subtleties relating to the choice of programming language, which do not affect the argument.
- [67] For non-periodic processes, 2^{C_μ} is the geometric average of the expected recurrence times of the causal states, i.e., still the average amount of time needed to return to a state after leaving it.
- [68] This does not mean they are never in equilibrium, in the economists' sense, however.